

ATTORNEY'S DOCKET NO. 06502.0374
Client Reference No. P5292

UNITED STATES PATENT APPLICATION

OF

GUY L. STEELE, JR.

FOR

TOTAL ORDER COMPARATOR UNIT FOR COMPARING VALUES
OF TWO FLOATING POINT OPERANDS

INCORPORATION BY REFERENCE

[001] Related U. S. Patent Application Serial No. not yet assigned, filed on even date herewith in the name of Guy L. Steele Jr. and entitled "Floating Point System That Represents Status Flag Information Floating Point Operand," assigned to the assignee of the present application, is hereby incorporated by reference.

[002] Related U. S. Patent Application Serial No. not yet assigned, filed on even date herewith in the name of Guy L. Steele Jr. and entitled "Comparator Unit For Comparing Values Of Floating point Operands," assigned to the assignee of the present publication, is hereby incorporated by reference.

FIELD OF THE INVENTION

[003] The invention relates generally to systems and methods for performing floating point operations, and more particularly to systems and methods for comparing a pair of floating point operands and generating a result that reflects a total ordering of the values of the operands.

BACKGROUND OF THE INVENTION

[004] Digital electronic devices, such as digital computers, calculators, and other devices, perform arithmetic calculations on values in integer, or "fixed point," format, in fractional, or "floating point" format, or both. IEEE Standard 754, (hereinafter "IEEE Std. 754" or "the Standard") published in 1985 by the Institute of Electrical and Electronic Engineers, and adopted by the American National Standards Institute (ANSI), defines several standard formats for expressing values in floating point format, and a number of aspects regarding behavior of computation in connection therewith. In accordance with IEEE Std.

754, a representation in floating point format comprises a plurality of binary digits, or "bits," having the structure:

$$se_{msb}...e_{lsb}f_{msb}...f_{lsb}$$

where bit "s" is a sign bit indicating whether the entire value is positive or negative, bits " $e_{msb}...e_{lsb}$ " comprise an exponent field represent the exponent "e" in unsigned binary biased format, and bits " $f_{msb}...f_{lsb}$ " comprise a fraction field that represents the fractional portion "f" in unsigned binary format ("msb" represents "most significant bit" and "lsb" represents "least significant bit"). The Standard defines two general formats, namely, a "single" format which comprises thirty-two bits, and a "double" format which comprises sixty-four bits. In the single format, there is one sign bit "s," eight bits " $e_7...e_0$ " comprising the exponent field and twenty-three bits " $f_{22}...f_0$ " comprising the fraction field. In the double format, there is one sign bit "s," eleven bits " $e_{10}...e_0$ " comprising the exponent field and fifty-two bits " $f_{51}...f_0$ " comprising the fraction field.

[005] As indicated above, the exponent field of the floating point representation " $e_{msb}...e_{lsb}$ " represents the exponent "E" in biased format. The biased format provides a mechanism by which the sign of the exponent is implicitly indicated. In particular, the bits " $e_{msb}...e_{lsb}$ " represent a binary encoded value "e" such that " $e=E+bias$." This allows the exponent E to extend from -126 to +127 in the eight-bit "single" format and from -1022 to +1023 in the eleven-bit "double" format, and provides for relatively easy manipulation of the exponents in multiplication and division operations, in which the exponents are added and subtracted, respectively.

[006] IEEE Std. 754 provides for several different formats with both the single and double formats, which are generally based on the bit patterns of the bits " $e_{msb}...e_{lsb}$ " comprising the exponent field and the bits $f_{msb}...f_{lsb}$ comprising the fraction field. As shown in prior art Fig. 4, if a number is represented such that all of the bits " $e_{msb}...e_{lsb}$ " of the exponent field are binary one's (that is, if the bits represent a binary-encoded value of "255" in the single format or "2047" in the double format) and all of the bits $f_{msb}...f_{lsb}$ of the fraction field are binary zeros, then the value of the number is positive infinity 310 or negative infinity 320, depending on the value of the sign bit "s." In particular, the value "v" is $v = (-1)^s \infty$ where " ∞ " represents the value "infinity." On the other hand, if all of the bits " $e_{msb}...e_{lsb}$ " of the exponent field are binary one's and if the bits $f_{msb}...f_{lsb}$ of the fraction field are not all zeros, then the value that is represented is deemed "not a number," abbreviated in the Standard by "NaN" 330.

[007] If a number has an exponent field in which the bits " $e_{msb}...e_{lsb}$ " are neither all binary ones nor all binary zeros (that is, if the bits represent a binary-encoded value between 1 and 254 in the single format or between 1 and 2046 in the double format), the number is said to be in a "normalized" format 360. For a number in the normalized format, the value represented by the number is $v = (-1)^s 2^{e-bias} (1.f_{msb}...f_{lsb})$, where "|" represents a concatenation operation. Effectively, in the normalized format, there is an implicit most significant digit having the value "one," so that the twenty-three digits in the fraction field of the single format, or the fifty-two digits in the fraction field of the double format, will effectively represent a value having twenty-four digits or fifty-three digits of precision, respectively, where the value is less than two, but not less than one.

[008] On the other hand, if a number has an exponent field in which the bits "e_{msb}...e_{lsb}" are all binary zeros, representing the binary-encoded value of "zero," and a fraction field in which the bits f_{msb}...f_{lsb} are not all zero, the number is said to be in a "de-normalized" format 370. For a number in the de-normalized format, the value represented by the number is $v = (-1)^s 2^{e-bias+1} (0.f_{msb}...f_{lsb})$. It will be appreciated that the range of values of numbers that can be expressed in the de-normalized format is disjoint from the range of values of numbers that can be expressed in the normalized format, for both the single and double formats. Finally, if a number has an exponent field in which the bits "e_{msb}...e_{lsb}" are all binary zeros, representing the binary-encoded value of "zero," and a fraction field in which the bits f_{msb}...f_{lsb} are all zero, the number has the value "zero." It will be appreciated that the value "zero" may be positive zero 340 or negative zero 350, depending on the value of the sign bit.

[009] Generally, floating point units to perform computations whose results conform to IEEE Std. 754 are designed to generate a result in response to a floating point instruction in three steps:

[010] (a) First, an approximation calculation step in which an approximation to the absolutely accurate mathematical result (assuming that the input operands represent the specific mathematical values as described by IEEE Std. 754) is calculated. This calculation is sufficiently precise as to allow this accurate mathematical result to be summarized by a sign bit, an exponent (typically represented using more bits than are used for an exponent in the standard floating point format), and some number "N" of bits of the presumed result fraction, plus a guard bit and a sticky bit. The value of the exponent will be such that the value of the

fraction generated in step (a) consists of a 1 before the binary point and a fraction after the binary point. The bits are calculated so as to obtain the same result as the following conceptual procedure (which is impossible under some circumstances to carry out in practice): calculate the mathematical result to an infinite number of bits of precision in binary scientific notation, and in such a way that there is no bit position in the significand such that all bits of lesser significance are 1-bits (this restriction avoids the ambiguity between, for example, 1.100000... and 1.011111... as representations of the value "one-and-one-half"); then let the N most significant bits of the infinite significand be used as the intermediate result significand, let the next bit of the infinite significand be the guard bit, and let the sticky bit be 0 if and only if ALL remaining bits of the infinite significand are 0-bits (in other words, the sticky bit is the logical OR of all remaining bits of the infinite fraction after the guard bit).

[011] (b) Second, a rounding step, in which the guard bit, the sticky bit, perhaps the sign bit, and perhaps some of the bits of the presumed significand generated in step (a) are used to decide whether to alter the result of step (a). For the rounding modes defined by IEEE Std. 754, this is a decision as to whether to increase the magnitude of the number represented by the presumed exponent and fraction generated in step (a). Increasing the magnitude of the number is done by adding 1 to the significand in its least significant bit position, as if the significand were a binary integer. It will be appreciated that, if the significand is all 1-bits, then magnitude of the number is "increased" by changing it to a high-order 1-bit followed by all 0-bits and adding 1 to the exponent. It will be further appreciated that,

[012] (i) if the result is a positive number, and

[013] (a) if the decision is made to increase, effectively the decision has been made to increase the value of the result, thereby rounding the result up (that is, towards positive infinity), but

[014] (b) if the decision is made not to increase, effectively the decision has been made to decrease the value of the result, thereby rounding the result down (that is, towards negative infinity);
and

[015] (ii) if the result is a negative number, and

[016] (a) if the decision is made to increase, effectively the decision has been made to decrease the value of the result, thereby rounding the result down, but

[017] (b) if the decision is made not to increase, effectively the decision has been made to increase the value of the result, thereby rounding the result up.

[018] (c) Finally, a packaging step, in which the result is packaged into a standard floating point format. This may involve substituting a special representation, such as the representation defined for infinity or NaN if an exceptional situation (such as overflow, underflow, or an invalid operation) was detected. Alternatively, this may involve removing the leading 1-bit (if any) of the fraction, because such leading 1-bits are implicit in the standard format. As another alternative, this may involve shifting the fraction in order to construct a denormalized number. As a specific example, we assume that this is the step that

forces the result to be a NaN if any input operand is a NaN. In this step, the decision is also made as to whether the result should be an infinity. It will be appreciated that, if the result is to be a NaN or infinity, any result from step (b) will be discarded and instead the appropriate representation will be provided as the result.

[019] In addition in the packaging step, floating point status information is generated, which is stored in a floating point status register. The floating point status information generated for a particular floating point operation typically includes indications, for example, as to whether

[020] (i) a particular operand is invalid for the operation to be performed ("invalid operation");

[021] (ii) if the operation to be performed is division and the divisor is zero ("division-by-zero");

[022] (iii) an overflow occurred during the operation ("overflow");

[023] (iv) an underflow occurred during the operation ("underflow"); and

[024] (v) the rounded result of the operation is not exact ("inexact").

[025] These conditions are typically represented by flags that are stored in the floating point status register, separate from the floating point operand. The floating point status information can be used to dynamically control the operations taken in response to certain instructions, such as conditional branch, conditional move, and conditional trap instructions that may be in the instruction stream subsequent to the floating point instruction. Also, the floating point status information may enable processing of a trap sequence, which

will interrupt the normal flow of program execution. In addition, the floating point status information may be used to affect certain ones of the functional unit control signals that control the rounding mode. IEEE Std. 754 also provides for accumulating floating point status information from, for example, results generated for a series or plurality of floating point operations.

[026] IEEE Std. 754 has brought relative harmony and stability to floating point computation and architectural design of floating point units. Moreover, its design was based on some important principles, and rests on a sensible mathematical semantics that eases the job of programmers and numerical analysts. It also supports the implementation of interval arithmetic, which may prove to be preferable to simple scalar arithmetic for many tasks. Nevertheless, IEEE Std. 754 has some serious drawbacks, including:

[027] (i) Modes, which include the rounding mode and may also include a traps enabled/disabled mode, flags representing the floating point status information that is stored in the floating point status register, and traps that are required to implement IEEE Std. 754 introduce implicit serialization between floating point instructions, and between floating point instructions and the instructions that read and write the flags and modes. Rounding modes can introduce implicit serialization because they are typically indicated as global state, although in some microprocessor architectures, the rounding mode is encoded as part of the instruction operation code, which will alleviate this problem to that extent. This implicit serialization makes the Standard difficult to implement coherently in today's superscalar and parallel microprocessor architectures without loss of performance.

[028] Implicit serialization occurs when programmers and designers try to avoid the problems caused if every floating point instruction uses, and can change, the same floating point status register. This can create problems if, for example, two instructions are executing in parallel in a microprocessor architectures featuring several CPUs running at once and both cause an update of the floating point status register. In such a case, the contents of the status register would likely be incorrect with respect to at least one of the instructions, because the other parallel instruction will have written over the original contents. Similar problems can occur in scalar processor architectures, in which several instructions are issued and processed at once. To solve this problem, programmers and designers make sure that floating point instructions that can affect the floating point status register are executed in serial fashion, one instruction completing before another begins.

[029] (ii) The implicit side effects of a procedure that can change the flags or modes can make it very difficult for compilers to perform optimizations on floating point code. To be safe, compilers for most languages assume that every procedure call is an optimization barrier.

[030] (iii) Global flags, such as those that signal certain modes, make it more difficult to do instruction scheduling where the best performance is provided by interleaving instructions of unrelated computations. Instructions from regions of code governed by different flag settings or different flag detection requirements cannot easily be interleaved when they must share a single set of global flag bits.

[031] (iv) Traps have been difficult to integrate efficiently into architectures and programming language designs for fine-grained control of algorithmic behavior.

[032] U.S. Patent Application Serial No. not yet assigned, filed on even date herewith in the name of Guy L. Steele Jr. and entitled "Floating Point System That Represents Point Status Flag Information With A Floating Point Operand," describes a floating point unit in which floating point status information is encoded in the representations of the results generated thereby. By encoding the floating point status information relating to a floating point operation in the result that is generated for the operation, the implicit serialization required by maintaining the floating point status information separate and apart therefrom can be obviated. The floating point unit includes a plurality of functional units, including an adder unit, a multiplier unit, a divider unit, a square root unit, a maximum/minimum unit, a comparator unit and a tester unit, all of which operate under control of functional unit control signals provided by a control unit. U.S. Patent Application Serial No. not yet assigned, filed on even date herewith in the name of Guy L. Steele, Jr. and entitled "Comparator Unit For Comparing Values Of Floating point Operands," describes a partial order comparator unit for use in connection with the floating point unit described in the aforementioned application that conforms to the requirements of IEEE Std. 754 that:

[033] (i) operands having values -0 and +0 (negative zero and positive zero) are to be considered equal to each other, and

[034] (ii) operands representing NaN values be considered "incomparable" (that is, not able to be compared) to any operands, including other NaN operands, regardless of the values represented by the bits comprising the sign, exponent, or fraction field(s) of the NaN operand(s) in the NaN format.

[035] While these requirements are not unreasonable in the context of, for example, floating point computation, it is desirable to relax them under some circumstances, such as, for example, when floating point values are used as sort keys. In that case, it will be desirable to use a consistent total order on all possible floating point values, including -0, +0 and NaN values. The instant application is directed to a total order comparison system that can be used in the floating point unit described in that application, as well as in other conventional floating point units.

[036] Systems and methods consistent with the principles of the present invention provide a total order comparison of operands, as the term "total order" is used in discrete mathematics. Total order refers to a relationship among a set of objects whereby for every pair of objects in the set, there is defined a relationship of "less than," "greater than," or "equal to." In a total order mathematical set, there are no "cannot compare" or "undefined" relationships. In contrast, in a partial order set, a "cannot compare" relationship may exist.

SUMMARY OF THE INVENTION

[037] Embodiments consistent with the principles of the present invention provide new and improved systems and methods for comparing a pair of floating point operands and generating a result that reflects a total ordering of the values of the operands.

[038] In one embodiment consistent with the principles of the present invention, a floating point total order comparator circuit for comparing a first floating point operand and a second floating point operand includes a first analysis circuit for determining a format of the first floating point operand based upon floating point status information encoded within the first floating point operand, a second analysis circuit for determining a format of the second

floating point operand based upon floating point status information encoded within the second floating point operand, and a result generator circuit coupled to the analysis circuits for producing a result indicating a total order comparative relationship between the first floating point operand and the second floating point operand based on the format of the first floating point operand and the format of the second floating point operand. Each floating point operand may include a fraction field, an exponent field, and a sign bit. The result may condition the outcome of a floating point instruction. The floating point total order comparator circuit may recognize several predetermined operand formats, such as not-a-number (NaN), infinity, normalized, denormalized, invalid operation, overflow, underflow, division by zero, positive zero, negative zero, exact, and inexact.

BRIEF DESCRIPTION OF THE DRAWINGS

[039] This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

[040] FIG. 1 is a functional block diagram of an exemplary total order comparator unit constructed in accordance with the principles of the invention;

[041] FIG. 2 depicts exemplary formats for representations of floating point values generated by the exemplary comparator unit depicted in FIG. 1;

[042] FIG. 3 is a functional block diagram of a second exemplary embodiment of a total order comparator unit constructed in accordance with the principles of the invention; and

[043] FIG. 4 depicts prior art formats for representations of floating point values.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

[044] FIG. 1 is a functional block diagram of a total order comparator unit 10 constructed in accordance with the principles of the invention. Generally, the total order comparator unit 10 receives two floating point operands and, in response during a comparison operation, generates signals that indicate whether the value of one operand is greater than the value of the other, the value of one operand is less than the value of the other, or the values of the two operands are equal. Thus, systems and methods consistent with the principles of the present invention provide a total order comparison of operands, as the term “total order” is used in discrete mathematics. Total order refers to a relationship among a set of objects whereby for every pair of objects in the set, there is defined a relationship of “less than,” “greater than,” or “equal to.” In a total order mathematical set, there are no “cannot compare” or “undefined” relationships. In contrast, in a partial order set, a “cannot compare” relationship may exist.

[045] As will be described below, in some cases, operands received by the exemplary total order comparator unit 10 may include floating point status information. In those cases, the floating point status information is encoded in and comprises part of the floating point representation of the operands, as well as the generated result. Since the floating point status information comprises part of the result, instead of being separate and apart from the operands (e.g., in a status register) as in prior art comparator units, the implicit serialization that is required by maintaining the floating point status information separate and apart from the result can be obviated. The total order comparator unit 10 facilitates the

comparison of values of the respective operands in a manner so as to accommodate the floating point status information that may be included therein.

[046] The exemplary total order comparator circuit 10 encodes the floating point status information in results that are generated in certain formats. FIG. 2 depicts exemplary formats of floating point operands that the total order comparator unit 10 may receive, and of results that it may generate. With reference to FIG. 2, seven exemplary formats are depicted, including a zero format 100, an underflow format 101, a denormalized format 102, a normalized non-zero format 103, an overflow format 104, an infinity format 105 and a not-a-number (NaN) format 106. The zero format 100 is used to represent the values "zero," or, more specifically, positive or negative zero, depending on the value of "s," the sign bit.

[047] The exemplary underflow format 101 provides a mechanism by which the floating point total order comparator circuit 10 can indicate that the result of a computation is an underflow. In the underflow format, the sign bit "s" indicates whether the result is positive or negative, the bits $e_{msb} \dots e_{lsb}$ of the exponent field are all binary zeros, and the bits $f_{msb} \dots f_{lsb+1}$ of the fraction field, except for the least significant bit, are all binary zeros. The least significant bit f_{lsb} of the fraction field is a binary one.

[048] The exemplary denormalized format 102 and normalized non-zero format 103 are used to represent finite non-zero floating point values substantially along the lines of that described above in connection with IEEE Std. 754. In both formats 102 and 103, the sign bit "s" indicates whether the result is positive or negative. The bits $e_{msb} \dots e_{lsb}$ of the exponent field of the denormalized format 102 are all binary zeros, whereas the bits $e_{msb} \dots e_{lsb}$ of the exponent field of the normalized non-zero format 103 are mixed ones and zeros,

except that the exponent field of the normalized non-zero format 103 will not have a pattern in which bits $e_{msb} \dots e_{lsb+1}$ are all binary ones and the least significant bit e_{lsb} zero and all of the bits $f_{msb} \dots f_{lsb}$ of the fraction field are all binary ones. In both formats 102 and 103, the bits $f_{msb} \dots f_{lsb}$ of the fraction field are not all binary zeros.

[049] The exemplary overflow format 104 provides a mechanism by which the floating point total order comparator circuit 10 can indicate that the result of a computation is an overflow. In the overflow format 104 the sign bit "s" indicates whether the result is positive or negative, the bits $e_{msb} \dots e_{lsb+1}$ of the exponent field are all binary ones, with the least significant bit e_{lsb} being zero. The bits $f_{msb} \dots f_{lsb}$ of the fraction field are all binary ones.

[050] The exemplary infinity format 105 provides a mechanism by which the floating point total order comparator circuit 10 can indicate that the result is infinite. In the infinity format 105, the sign bit "s" indicates whether the result is positive or negative, the bits $e_{msb} \dots e_{lsb}$ of the exponent field are all binary ones, and the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field are all binary zeros. The five least significant bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field are flags, which will be described below.

[051] The exemplary NaN (not-a-number) format 106 provides a mechanism by which the floating point total order comparator circuit 10 can indicate that the result is not a number. In the NaN format, the sign bit "s" can be any value, the bits $e_{msb} \dots e_{lsb}$ of the exponent field are all binary ones, and the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field are not all binary zeros. The five least significant bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field are flags, which will be described below.

[052] As noted above, in values represented in the infinity format 105 and the NaN format 106, the five low order bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field are flags. In the formats used with the floating point total order comparator circuit 10, the five flags include the flags that are defined by IEEE Std. 754, including an invalid operation flag "n," an overflow flag "o," an underflow flag "u," a division-by-zero flag "z" and an inexact flag "x." For example, a value in the NaN format 106 in which both the overflow flag "o" and the division-by-zero flag "z" are set indicates that the value represents a result of a computation that involved an overflow (this from the overflow flag "o"), as well as an attempt to divide by zero (this from the division-by-zero flag "z"). It should be noted that the flags provide the same status information as provided by, for example, information stored in a floating point status register in a prior art floating point unit. However, since the information is provided as part of the result and stored therewith in registers in which the result is ultimately stored, multiple instructions can be contemporaneously executed, as the floating point status information that may be generated during execution of one instruction, when stored, will not over-write previously-stored floating point status information generated during execution of another instruction. In addition, in one embodiment, values in the other formats can be indicated as being inexact based in part on the least-significant bit f_{lsb} of their fraction fields, and in that embodiment that bit operates as an inexact flag. The value will be indicated as being inexact if the bit f_{lsb} has the value "one," and otherwise exact.

[053] One of ordinary skill in the art will recognize that the other formats and bit patterns could be used to represent the floating point operand formats without departing from the principles of the present invention. One of ordinary skill in the art will also recognize that the floating point status information contained in the operands could easily be

represented by other bit combinations (not shown) without departing from the principles of the present invention. For example, more or fewer bits could be used, a subset or superset of the exemplary status bits could be used, or the most significant bits of an operand (or some other subset of bits) could be used to indicate the floating point status information, instead of the least significant bits illustrated.

[054] With this background the structure and operation of the exemplary total order comparator unit 10 will be described in connection with FIG. 1. With reference to FIG. 1, exemplary total order comparator unit 10, consistent with the principles of the present invention, includes two operand buffers 11A and 11B, respective operand analysis circuits 12A and 12B, and a result generator circuit 13. The operand buffers 11A and 11B receive and store respective operands that may be received from, for example, a set of registers (not shown) in a conventional manner. Each operand analysis circuit 12A, 12B analyzes the operand in the respective operand buffer 11A, 11B and generates signals providing information relating to the respective operands, which signals are provided to the result generator circuit 13. The signals provided by the operand analysis circuit 12A, 12B essentially provide information as to the type of the respective operand, in particular indicating whether the operand is in the infinity format 105 or the NaN format 106. The result generator 13 receives the signals from the operand analysis circuits 12A, 12B and signals representative of the operands in the operand buffers 11A and 11B and generates three result signals, namely an A LT B (operand A is less than operand B) signal, an A GT B (operand A is greater than operand B) signal, and an A EQ B (operand A is equal to operand B) signal, precisely one of which will be asserted to indicate the comparative relationship between the values of the operands in the operand buffers 11A and 11B.

[055] Before proceeding to a detailed description of the operand analysis circuits 12A, 12B and result generator circuit 13, it would be helpful to describe the exemplary result that will be generated by the exemplary total order comparator unit 10 based on the operands in operand buffers 11A, 11B. Generally, and consistent with one embodiment of the present invention:

[056] (a) if the sign bits, bits $e_{msb} \dots e_{lsb}$ comprising the exponent fields, and bits $f_{msb} \dots f_{lsb}$ of the operands are the same, regardless of the formats of the respective operands and regardless of whether the bits represent floating point status flags, the operands will be considered equal;

[057] (b) an operand having the value +0 (that is, an operand in the zero format 100 with the sign bit clear) is considered as having a larger value than an operand having the value -0 (that is, an operand in the zero format 100 with the sign bit set);

[058] (c) if the operands are in the zero format 100, underflow format 101, denormalized format 102, normalized non-zero format 103, overflow format 104, or infinity format 105, the result will reflect their respective numerical values, with the recognition that:

[059] (I) the magnitude (that is, the absolute value) of an operand in the zero format 100 is less than the magnitude of an operand in the underflow format 101,

[060] (II) the magnitude of an operand in the underflow format 101 is less than the magnitude of an operand in the denormalized format 102,

[061] (III) the magnitude of an operand in the denormalized format 102 is less than the magnitude of an operand in the normalized non-zero format 103,

[062] (IV) the magnitude of an operand in the normalized non-zero format 103 is less than the magnitude of an operand in the overflow format 104, and

[063] (V) the magnitude of an operand in the overflow format 104 is less than the magnitude of an operand in the infinity format 105, regardless of the condition of the bits $f_{lsb+4} \dots f_{lsb}$ comprising the floating point status flags of the operand in the infinity format;

[064] (d) if both operands are in the infinity format 105 and their sign bits are the same, in considering the binary-encoded values of the bits $f_{lsb+4} \dots f_{lsb}$ of the fraction fields of the respective operands, which represent the floating point status flags, the operand with the greater binary-encoded value will be considered to be greater than the other operand;

[065] (e) any operand in the NaN format 106, regardless of sign, is greater than any operand in any other (non-NaN) format; and

[066] (f) if both operands are in the NaN format, then

[067] (I) if the binary-encoded values of bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the two operands (that is, the portions of the fraction fields except for bits $f_{lsb+4} \dots f_{lsb}$ comprising the floating point status flags) are the same, and if the signs of the two operands differ, the positive operand will be considered to be the greater;

[068] (II) if the binary-encoded values of bits $f_{msb} \dots f_{lsb+5}$ of the fraction fields of the two operands are not the same, the operand for which the binary-encoded value of bits $f_{msb} \dots f_{lsb+5}$ of the fraction field is larger is considered to be the greater (regardless of the sign bits); and

[069] (III) if the binary-encoded values of bits $f_{msb} \dots f_{lsb+5}$ of the fraction fields of the two operands are the same and the signs are the same, but the binary-encoded values of bits $f_{lsb+4} \dots f_{lsb}$ of the fraction fields of the two operands are not the same, the operand for which the binary-encoded value of bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field is larger is considered to be the greater.

[070] One of ordinary skill in the art will recognize that the exemplary total order comparator circuit 10 could be easily implemented in alternative embodiments to generate a different result for given operands in operand buffers 11A, 11B, without departing from the principles of the present invention. The circuit could be easily implemented to reflect a different embodiment of the “total order” than the one illustrated by (a) through (f). For example, an alternative embodiment of the total order may be one in which an operand in the negative NaN format 106 (i.e., with the sign bit set) is less than any positive operand in any other format. One of ordinary skill could easily use a new total order result function to meet specific design needs and implement it by changing the particular implementation of the exemplary circuits shown in FIG. 1, or by designing new circuits in accordance with the principles of the present invention.

[071] With this background, referring to FIG. 1, each operand analysis circuit 12A, 12B analyzes the operand in the respective buffer 11A, 11B and generates signals providing information relating thereto, which signals are provided to the result generator 13. In the illustrated embodiment, each exemplary operand analysis circuit 12A, 12B comprises a number of elements, including:

[072] (i) a comparator 20A, 20B that generates an asserted signal if the bits $e_{msb} \dots e_{lsb}$ of the exponent field of the operand in respective buffer 11A, 11B are all binary ones, which will be the case if the operand is in the infinity format 105 or the NaN format 106;

[073] (ii) a comparator 21A, 21B that generates an asserted signal if the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in respective buffer 11A, 11B are all binary zeros, which will be the case if, inter alia, the operand is in the infinity format 105 and will not be the case if, inter alia, the operand is in the NaN format 106;

[074] (iii) an AND gate 22A, 22B that generates an asserted A INF or B INF signal if the comparator 20A, 20B and the comparator 21A, 21B are both generating asserted signals; it will be appreciated that AND gate 22A, 22B will be generating an asserted signal if the operand in operand buffer 11A, 11B is in the infinity format 105; and

[075] (iv) a NAND gate 23A, 23B that generates an asserted A NaN or B NaN signal if the comparator 20A, 20B is generating an asserted signal and the comparator 21A, 21B is generating a negated signal; it will be appreciated that AND gate 23A, 23B will be generating an asserted signal if the operand in operand buffer 11A, 11B is in the NaN format 106.

[076] As noted above, the exemplary result generator 13 receives the signals from the operand analysis circuits 12A, 12B and signals representative of the operands in the operand buffers 11A and 11B. Result generator circuit 13 generates three result signals, namely an A LT B (operand A is less than operand B) signal, an A GT B (operand A is greater than operand B) signal, and an A EQ B (operand A is equal to operand B) signal, precisely one of which will be asserted to indicate the total order comparative relationship between the values of the operands in the operand buffers 11A and 11B. In the illustrated embodiment, the result generator 13 includes elements (v) through (xxvi) as follows:

[077] (v) a comparator 30 that generates an asserted signal if the binary-encoded values of the bits $e_{msb} \dots e_{lsb} f_{msb} \dots f_{lsb}$ comprising the exponent and fraction fields of the operands in the operand buffers 11A and 11B are the same;

[078] (vi) an XOR gate 31 that generates a negated signal if both an A NEG signal representative of the sign bit in operand buffer 11A and a B NEG signal representative of the sign bit in operand buffer 11B are both asserted, or both negated, which will be the case if the operands in both operand buffers 11A and 11B are negative or if the operands in both operand buffers are positive;

[079] (vii) a NAND gate 32 that generates an asserted A EQ B (operand A equals operand B) signal if the comparator 30 is generating an asserted signal and the XOR gate 31 is generating a negated signal; it will be appreciated that the NAND gate 32 will be generating an asserted A EQ B signal if the operands in the two operand buffers 11A and 11B have the same value (reference item (a) above);

[080] (viii) a comparator 33 that generates an asserted signal if the binary-encoded value of the bits $e_{msb} \dots e_{lsb}$ and $f_{msb} \dots f_{lsb}$ comprising the exponent and fraction fields of the operand in operand buffer 11A is less than the binary-encoded value of the value of the same bits of the operand in operand buffer 11B; it will be appreciated that comparator 33 will be generating an asserted signal if the magnitude of the operand in operand buffer 11B is greater than the magnitude of the operand in operand buffer 11A;

[081] (ix) an AND gate 34 that generates an asserted signal if both the A NEG and the B NEG signals are asserted; it will be appreciated that the AND gate 34 will be generating an asserted signal if the operands in both operand buffers are negative;

[082] (x) a NAND gate 35 that generates an asserted signal if the A NaN signal from NAND gate 23A is negated, the signal generated by comparator 20B is negated, the signal generated by comparator 33 is negated, and the signal generated by AND gate 34 is asserted; it will be appreciated that NAND gate 35 will be generating an asserted signal if the operands in both operand buffers 11A and 11B are negative, the magnitude of the operand in operand buffer 11B is less than or equal to the magnitude of the operand in operand buffer 11A, the operand in operand buffer 11A is not in the NaN format 106, and the operand in operand buffer 11B is not in either the infinity format 105 or the NaN format 106, in which case the value of the operand in operand buffer 11A will be less than or equal to the value of the operand in operand buffer 11B (reference items (b) and (c) above);

[083] (xi) a NAND gate 36 that generates an asserted signal if both the A NEG and the B NEG signals are negated; it will be appreciated that the AND gate 34 will be generating an asserted signal if the operands in both operand buffers are positive;

[084] (xii) a NAND gate 37 that generates an asserted signal if the comparator 20A is generating a negated signal, the B NaN signal from NAND gate 23B is negated, and the signals generated by comparator 33 and NAND gate 36 are both asserted; it will be appreciated that NAND gate 35 will be generating an asserted signal if the operands in both operand buffers 11A, 11B are positive, the magnitude of the operand in operand buffer 11B is greater than the magnitude of the operand in operand buffer 11A, the operand in operand buffer 11B is not in the NaN format 106, and the operand in operand buffer 11A is not in either the infinity format 105 or the NaN format 106, in which case the value of the operand in operand buffer 11A will be less than or equal to the value of the operand in operand buffer 11B (reference items (b) and (c) above);

[085] (xiii) a NAND gate 40 that generates an asserted signal if the A NaN signal from NAND gate 23A is negated and the B NaN signal from NAND gate 23B is asserted; it will be appreciated that the NAND gate 40 will be generating an asserted signal if the operand in operand buffer 11A is not in the NaN format 106 and the operand in operand buffer 11B is in the NaN format 106 (reference item (e) above);

[086] (xiv) a comparator 41 that generates an asserted signal if the binary-encoded value represented by signals representative of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11A is less than the binary-encoded value represented by signals representative of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11B;

[087] (xv) a comparator 42 that generates an asserted signal if the binary-encoded value represented by signals representative of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the

operand in operand buffer 11A is equal to the binary-encoded value represented by signals representative of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11B;

[088] (xvi) a comparator 43 that generates an asserted signal if the binary-encoded value represented by signals representative of the bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field of the operand in operand buffer 11A is less than the binary-encoded value represented by signals representative of the bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field of the operand in operand buffer 11B; it will be appreciated that, if the operands in operand buffers 11A and 11B are in the infinity format 105 or the NaN format 106, the bits $f_{lsb+4} \dots f_{lsb}$ correspond to the flags "n," "o," "u," "z" and "x," respectively;

[089] (xvii) a NAND gate 44 that generates an asserted A-B+ (operand A negative, operand B positive) signal if the A NEG signal is asserted and the B NEG signal is negated; it will be appreciated that the NAND gate 44 will be generating an asserted signal if the operand in operand buffer 11A is negative and the operand in operand buffer 11B is positive;

[090] (xviii) an AND gate 45 that generates an asserted signal if both the comparator 42 and the NAND gate 44 are generating asserted signals; it will be appreciated that the AND gate 45 will be generating an asserted signal if the binary-encoded value of bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11A is equal to the binary-encoded value of the corresponding bits of the operand in operand buffer 11B, and the operand in operand buffer 11A is negative while the operand in operand buffer 11B is positive;

[091] (xix) a NAND gate 46 that generates an asserted signal if the comparator 42 and comparator 43 are generating asserted signals and the XOR gate 31 is generating a negated SA NEQ SB (sign operand A not equal to sign operand B) signal; it will be appreciated that the NAND gate 46 will be generating an asserted signal if the operands in the operand buffers 11A and 11B have the same sign, the binary-encoded values of the bits $f_{msb} \dots f_{lsb+5}$ of the operands in the operand buffers 11A and 11B are the same, and the binary-encoded value of the bits $f_{lsb+4} \dots f_{lsb}$ of the operand in operand buffer 11A is less than the binary-encoded value of the corresponding bits of the operand in operand buffer 11B;

[092] (xx) an OR gate 47 that generates an asserted signal if any of comparator 41, AND gate 45 or NAND gate 46 is generating an asserted signal;

[093] (xxi) an AND gate 48 that generates an asserted signal if NAND gates 23A and 23B are generating respective asserted A NaN and B NaN signals and OR gate 47 is generating an asserted signal; it will be appreciated that AND gate 48 will be generating an asserted signal if the operands in both operand buffers 11A and 11B are in the NaN format 106 and either:

[094] (A) the binary-encoded value of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11A is less than the binary-encoded value of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11B (reference comparator 41 and item (f)(I) above); or

[095] (B) the binary-encoded value of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11A equals the binary-encoded value of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand

buffer 11B, and the operand in operand buffer 11A is negative while the operand in operand buffer 11B is positive (reference AND gate 45 and item (f)(II) above); or

- [096] (C) the signs of the operands in both operand buffers 11A and 11B are the same, and the binary-encoded value of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11A equals the binary-encoded value of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11B, but the binary-encoded value of the bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field of the operand in operand buffer 11A is less than the binary-encoded value of the bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field of the operand in operand buffer 11B (reference NAND gate 46 and item (f)(III) above);

[097] (xxii) a NAND gate 50 that generates an asserted signal if the NAND gates 23A and 23B are generating negated A NaN and B NaN signals and if the NAND gate 44 is generating an asserted A- B+ (operand A negative, operand B positive) signal; it will be appreciated that AND gate 50 will be generating an asserted signal if neither operand is in the NaN format 106 and the operand in operand buffer 11A is negative and the operand in operand buffer 11B is positive, which will ensure that all other operands, except operands in the NaN format 106, that are negative will be less than all operands, other than operands in the NaN format 106, that are positive (reference items (b) and (c) above);

[098] (xxiii) a NAND gate 51 that generates an asserted signal if the AND gates 22A and 22B are generating asserted A INF and B INF signals, XOR gate 31 is generating a

negated SA NEQ SB (sign operand A not equal to sign operand B) signal, and comparator 43 is generating an asserted signal; it will be appreciated that NAND gate 51 will be generating an asserted signal if the operands in both operand buffers 11A and 11B are in the infinity format 105, the signs of the operands are the same, and the binary-encoded value of bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field of the operand in operand buffer 11A (which represents the operand's floating point status flags) is less than the binary-encoded value of bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field of the operand in operand buffer 11B (reference item (d) above);

[099] (xxiv) an OR gate 52 that generates an asserted signal if any of NAND gates 35, 37, 40, 50 or 51 or AND gate 48 is generating an asserted signal; it will be appreciated that the OR gate 52 will be generating an asserted signal if:

[0100] (A) the operands in both operand buffers 11A, 11B are negative, the magnitude of the operand in operand buffer 11B is less than or equal to the magnitude of the operand in operand buffer 11A, the operand in operand buffer 11A is not in the NaN format 106 and the operand in operand buffer 11B is not in either the infinity format 105 or the NaN format 106, in which case the value of the operand in operand buffer 11A will be less than or equal to the value of the operand in operand buffer 11B (reference NAND gate 35, items (a) and (b) above);

[0101] (B) the operands in both operand buffers 11A, 11B are positive, the magnitude of the operand in operand buffer 11B is greater than the magnitude of the operand in operand buffer 11A, the operand in operand buffer 11B is not in the NaN format 106, and the operand in

operand buffer 11A is not in either the infinity format 105 or the NaN format 106, in which case the value of the operand in operand buffer 11A will be less than or equal to the value of the operand in operand buffer 11B (reference NAND gate 37, items (a) and (b) above);

[0102] (C) the operand in operand buffer 11A is not in the NaN format 106 and the operand in operand buffer 11B is in the NaN format 106 (reference NAND gate 40, item (e) above);

[0103] (D) the operands in both operand buffers 11A and 11B are in the NaN format 106 (reference AND gate 48) and either:

[0104] (1) the binary-encoded value of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11A is less than the binary-encoded value of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11B (reference comparator 41, item (f)(I) above);

[0105] (2) the binary-encoded value of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11A equals the binary-encoded value of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11B, but the operand in operand buffer 11A is negative while the operand in operand buffer 11B is positive (reference AND gate 45 above, item (f)(II) above);
or

[0106] (3) the signs of the operands in both operand buffers 11A and 11B are the same, and the binary-encoded value of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11A equals the binary-encoded value of the bits $f_{msb} \dots f_{lsb+5}$ of the fraction field of the operand in operand buffer 11B, but the binary-encoded value of the bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field of the operand in operand buffer 11A is less than the binary-encoded value of the bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field of the operand in operand buffer 11B (reference AND gate 46, item (f)(III) above);

[0107] (E) neither operand is in the NaN format 106 and the operand in operand buffer 11A is negative and the operand in operand buffer 11B is positive, which will ensure that all operands (other than those in the NaN format 106) that are negative will be less than all operands (other than operands in NaN format 106) that are positive (reference NAND gate 50, items (b) and (c) above); or

[0108] (F) the operands in both operand buffers 11A and 11B are in the infinity format 105, and the signs of the operands are the same, and the binary-encoded value of bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field of the operand in operand buffer 11A (which represents the operand's flags) is less than the binary-encoded value of bits $f_{lsb+4} \dots f_{lsb}$ of the fraction field of the

operand in operand buffer 11B (reference NAND gate 51, item (d) above).

[0109] (xxv) a NAND gate 53 that generates an asserted A LT B (operand A less than operand B) signal if the OR gate 52 is generating an asserted signal and the NAND gate 32 is generating a negated signal; it will be appreciated that the NAND gate 53 will be generating an asserted signal whenever the OR gate 52 is generating an asserted signal as described in connection with item (xxiv) above, except when NAND gate 35 is generating an asserted signal because the magnitude of the operand in operand buffer 11A equals the magnitude of the operand in operand buffer 11B (reference item (xxiv)(A) above); and

[0110] (xxvi) a NAND gate 54 that generates an asserted A GT B (operand A greater than operand B) signal if the OR gate 52 and the NAND gate 32 are both generating negated signals; it will be appreciated that the NAND gate 54 will be generating an asserted signal whenever the OR gate 52 is generating a negated signal as described in connection with item (xxiv) above, except when NAND gate 32 is generating an asserted signal because the magnitude of the operand in operand buffer 11A equals the magnitude of the operand in operand buffer 11B (reference item (xxiv)(A) above).

[0111] As noted above, the exemplary total order comparator unit 10 can provide the A LT B, A GT B, and A EQ B signals to a control unit for controlling a floating point unit that includes the total order comparator unit 10, as described in the aforementioned related application.

[0112] Systems consistent with the principles of the invention provides a number of features. In particular, such systems provide a total order comparator unit 10 that provides

comparisons that indicate a consistent total order on all possible floating point values, including -0, +0, and NaN values for operands for which floating point status information forms part of respective ones of the operands.

[0113] The exemplary total order comparator unit 10 described above in connection with FIG. 1 is designed to accommodate operands wherein the floating point status flags form part of the operands in several of the formats. FIG. 3 depicts another exemplary total order comparator unit 110 that can be used in floating point units that conform to existing IEEE STD. 754 operands, in which floating point status information does not form part of the respective operands, but instead is maintained separate and apart from the operands. With reference to the embodiment illustrated in FIG. 3, exemplary total order comparator unit 110 includes two operand buffers 111A and 111B, respective operand analysis circuits 112A and 112B, and a result generator circuit 113. The operand buffers 111A and 111B receive and store respective operands that may be received from, for example, a set of registers (not shown) in a conventional manner. Each operand analysis circuit 112A, 112B analyzes the operand in the respective operand buffer 111A, 111B and generates signals providing information relating to the respective operands, which signals are provided to the result generator 113. The signals provided by the operand analysis circuit 112A, 112B essentially provide information as to the type of the respective operand, in particular indicating whether the operand is an NaN format. The result generator 113 receives the signals from the operand analysis circuits 112A, 112B and signals representative of the operands in the operand buffers 111A and 111B. Result generator circuit 113 generates three result signals, namely an A LT B (operand A is less than operand B) signal, an A GT B (operand A is greater than operand B) signal, and an A EQ B (operand A is equal to operand B) signal, precisely one of which

will be asserted to indicate the total order comparative relationship between the values of the operands in the operand buffers 111A and 111B.

[0114] Before proceeding to a detailed description of the exemplary operand analysis circuits 112A, 112B and result generator 113, it would be helpful to describe the result that will be generated by the exemplary total order comparator unit 110 based on the operands in operand buffers 111A, 111B. Generally, and consistent with one embodiment of the present invention:

[0115] (a) if the sign bits, bits $e_{msb} \dots e_{lsb}$ comprising the exponent fields and $f_{msb} \dots f_{lsb}$ of the operands are the same, regardless of the values of the respective operands, they will be considered equal;

[0116] (b) an operand having the value +0 is considered as having a larger value than an operand having the value -0;

[0117] (c) if neither operand is a NaN, the result will reflect their respective numerical values;

[0118] (d) any operand in the NaN format 106, regardless of sign, is greater than any operand in any other format; and

[0119] (e) if both operands are NaN's, the result will reflect the binary-encoded values of the bits comprising the respective operands,

[0120] (I) excluding the signs of the respective operands if the magnitudes are not equal, but

[0121] (II) including the signs of the respective operands if the magnitudes are equal.

[0122] With this background, referencing FIG. 3, each exemplary operand analysis circuit 112A, 112B analyzes the operand in the respective buffer 111A, 111B and generates signals providing information relating thereto, which signals are provided to the result generator circuit 113. In the illustrated embodiment, each operand analysis circuit 112A, 112B, comprises a number of elements, including:

[0123] (i) a comparator 120A, 120B that generates an asserted signal if the bits $e_{msb} \dots e_{lsb}$ of the exponent field of the operand in respective buffer 111A, 111B are all binary ones, which will be the case if the operand is, inter alia, a NaN;

[0124] (ii) a comparator 121A, 121B that generates an asserted signal if the bit $f_{msb} \dots f_{lsb}$ of the fraction field of the operand in respective buffer 11A, 11B are all binary zeros, which will not be the case if the operand is, inter alia, a NaN;

[0125] (iii) a NAND gate 123A, 123B that generates an asserted A NaN or B NaN signal if the comparator 120A, 120B is generating an asserted signal and the comparator 121A, 121B is generating a negated signal; it will be appreciated that NAND gate 123A, 123B will be generating an asserted signal if the operand in operand buffer 111A, 111B is a NaN.

[0126] As noted above, the result generator 113 receives the signals from the operand analysis circuits 112A, 112B and signals representative of the operands in the operand buffers 111A and 111B and generates three result signals, namely an A LT B (operand A is less than operand B) signal, an A GT B (operand A is greater than operand B)

signal, and an A EQ B (operand A is equal to operand B) signal, precisely one of which will be asserted to indicate the total order comparative relationship between the values of the operands in the operand buffers 111A and 111B. In the illustrated embodiment, the result generator 113 includes:

[0127] (iv) a comparator 130 that generates an asserted signal if the binary-encoded values of the bits $e_{msb} \dots e_{lsb} f_{msb} \dots f_{lsb}$ comprising the exponent and fraction fields of the operands in the operand buffers 11A and 11B are the same;

[0128] (v) an XOR gate 131 that generates a negated signal if an A NEG signal representative of the sign bit in operand buffer 111A and a B NEG signal representative of the sign bit in operand buffer 111B are both asserted, or both negated, which will be the case if the operands in both operand buffers 111A and 111B are negative or if the operands in both operand buffers are positive;

[0129] (vi) a NAND gate 132 that generates an asserted A EQ B (operand A equals operand B) signal if the comparator 130 is generating an asserted signal and the XOR gate 131 is generating a negated signal; it will be appreciated that the NAND gate 132 will be generating an asserted A EQ B signal if the operands in the two operand buffers 111A and 111B have the same value (reference item (a) directly above);

[0130] (vii) a comparator 133 that generates an asserted signal if the binary-encoded value of the bits $e_{msb} \dots e_{lsb}$ and $f_{msb} \dots f_{lsb}$ comprising the exponent and fraction fields of the operand in operand buffer 111A is less than the binary-encoded value of the value of the same bits of the operand in operand buffer 111B; it will be appreciated that comparator

133 will be generating an asserted signal if the magnitude of the operand in operand buffer 111B is greater than the operand in operand buffer 111A;

[0131] (viii) an XOR gate 134 that generates an asserted signal if a B NEG signal representative of the sign bit of the operand in operand buffer 111B is negated and the comparator 133 is generating an asserted signal, or if the B NEG signal is asserted and the comparator 133 is generating a negated signal; it will be appreciated that the XOR gate 134 will be generating an asserted signal if:

[0132] (A) the operand in operand buffer 111B is negative and the magnitude of the operand in operand buffer 111B is less than the magnitude of the operand in operand buffer 111A, in which case the value of the operand in operand buffer 111A is less than the value of the operand in operand buffer 111B if the operand in operand buffer 111A is negative, or

[0133] (B) the operand in operand buffer 111B is positive and the magnitude of the operand in operand buffer 111A is less than the magnitude of the operand in operand buffer 111B, in which case the value of the operand in operand buffer 111A is less than the value of the operand in operand buffer 111B,

[0134] note that representations of operands that are NaNs are considered as if they represented numerical values;

[0135] (ix) a NAND gate 136 that generates an asserted signal if the XOR gate 131 is generating a negated signal and the XOR gate 134 is generating an asserted signal; it will

be appreciated that the NAND gate 136 will be generating an asserted signal if both operands are of the same sign and the value of the operand in operand buffer 111A is less than the value of the operand in operand buffer 111B, with representations of operands that are NaN's being considered to represent numerical values (reference item (c) directly above);

[0136] (x) a NAND gate 137 that generates an asserted signal if the A NEG signal is asserted and the B NEG signal is negated; it will be appreciated that the NAND gate 137 will be generating an asserted signal if the operand in operand buffer 111A is negative and the operand in operand buffer 111B is positive (reference items (b) and (c) directly above);

[0137] (xi) an OR gate 138 that generates an asserted signal if either NAND gate 136 or NAND gate 137 is generating an asserted signal; it will be appreciated that OR gate 138 will be generating an asserted signal either if the operand in operand buffer 111A is negative and the operand in operand buffer 111B is positive (reference NAND gate 137), or if the operands in the operand buffers have the same sign and the operand in operand buffer 111A is less than the operand in operand buffer 111B (reference NAND gate 136);

[0138] (xii) a NAND gate 139 that generates an asserted signal if OR gate 138 is generating an asserted signal and both the A NaN signal from NAND gate 123A and the B NaN signal from NAND gate 123B are negated; it will be appreciated that the NAND gate 139 will be generating an asserted signal if neither operand is a NaN and the operand in operand buffer 111A is less than the operand in operand buffer 111B (reference item (c) directly above);

[0139] (xiii) an AND gate 140 that generates an asserted signal if the comparator 130 and NAND gate 137 are both generating an asserted signal; it will be appreciated that the

AND gate 140 will be generating an asserted signal if the operands in the operand buffers 111A and 111B have the same magnitude, and the operand in operand buffer 111B is positive and the operand in operand buffer 111A is negative (reference item (e)(II) directly above);

[0140] (xiv) an OR gate 141 that generates an asserted signal if either the comparator 133 is generating an asserted signal or the AND gate 140 is generating an asserted signal; it will be appreciated that the OR gate 141 will be generating an asserted signal if:

[0141] (I) the operands in the operand buffers 111A and 111B have the same magnitude, but the operand in operand buffer 111B is positive and the operand in operand buffer 111A is negative (reference item (e)(II) directly above), or

[0142] (II) the magnitude of the operand in operand buffer 111A is less than the magnitude of the operand in operand buffer 111B (reference item (e)(I) directly above);

[0143] (xv) an AND gate 142 that generates an asserted signal if the OR gate 141 is generating an asserted signal, and if AND gate 123A and AND gate 123B are both generating asserted signals indicating that both operands are NaNs;

[0144] (xvi) a NAND gate 143 that generates an asserted signal if the NAND gate 123A is generating a negated signal and the NAND gate 123B is generating an asserted signal; it will be appreciated that the NAND gate 143 will be generating an asserted signal if the operand in operand buffer 111A is not a NaN and the operand in operand buffer 111B is a NaN (reference item (d) directly above);

[0145] (xvii) an OR gate 152 that generates an asserted signal if any of gates 139, 142 or 143 is generating an asserted signal; it will be appreciated that the OR gate 152 will be generating an asserted signal if:

[0146] (I) neither operand is a NaN and the operand in operand buffer 111A is less than the operand in operand buffer 111B (reference NAND gate 139 and items (b) and (c) directly above), or

[0147] (II) both operands are NaNs and

[0148] (A) the operands in the operand buffers 111A and 111B have the same magnitude, and the operand in operand buffer 111B is positive and the operand in operand buffer 111A is negative (reference item (e)(II) directly above), or

[0149] (B) the magnitude of the operand in operand buffer 111A is less than the magnitude of the operand in operand buffer 111B (reference AND gate 142 item (e)(I) directly above), or

[0150] (III) the operand in operand buffer 111A is not a NaN and the operand in operand buffer 111B is a NaN (reference NAND gate 143 and item (d) directly above);

[0151] (xviii) a NAND gate 153 that generates an asserted A LT B (operand A less than operand B) signal if the OR gate 152 is generating an asserted signal and the NAND gate 132 is generating a negated signal;

[0152] (xix) a NAND gate 154 that generates an asserted A GT B (operand A greater than operand B) signal if the OR gate 152 and NAND gate 132 are both generating negated signals; the NAND gate 154 makes use of the signal from NAND gate 132 since the signal from OR gate 152 can be negated if the sign bits, bits $e_{msb} \dots e_{lsb}$ comprising the exponent fields and $f_{msb} \dots f_{lsb}$ of the operands are the same (reference item (a) directly above).

[0153] One of ordinary skill in the art will recognize that the circuits implementing the comparative relationships between the operands specified in the preceding descriptions to produce the result can easily be altered to implement different comparative relationships and generate different results without departing from the principles of the present invention. For example, the relationships in an alternative embodiment may be defined and implemented so that the generated result is A EQ B if operand A is in the positive zero format and operand B is in the negative zero format, instead of generating an A GT B result as is done by the exemplary circuits of the illustrated embodiment.

[0154] It will be further appreciated that a system in accordance with the invention can be constructed in whole or in part from special purpose hardware or a general purpose computer system, or any combination thereof, any portion of which may be controlled by a suitable program. Any program may in whole or in part comprise part of or be stored on the system in a conventional manner, or it may in whole or in part be provided to the system over a network or other mechanism for transferring information in a conventional manner. In addition, it will be appreciated that the system may be operated and/or otherwise controlled by means of information provided by an operator using operator input elements (not shown), which may be connected directly to the system or which may transfer the information to the

system over a network or other mechanism for transferring information in a conventional manner.

[0155] Those skilled in the art will appreciate that the invention may be practiced in an electrical circuit comprising discrete electronic elements, packaged or integrated electronic chips containing logic gates, a circuit utilizing a microprocessor, or on a single chip containing electronic elements or microprocessors. It may also be provided using other technologies capable of performing logical operations such as, for example, AND, OR, and NOT, including but not limited to mechanical, optical, fluidic, and quantum technologies. In addition, the invention may be practiced within a general purpose computer or in any other circuits or systems as are known by those skilled in the art.

[0156] The foregoing description has been limited to specific embodiments of this invention. It will be apparent, however, that various variations and modifications may be made to the invention, with the attainment of some or all of the advantages of the invention. It is the object of the appended claims to cover these and such other variations and modifications as come within the true spirit and scope of the invention.